Defense Technical Information Center Compilation Part Notice

ADP012342

TITLE: Issues of Dynamic Coalition Formation Among Rational Agents

DISTRIBUTION: Approved for public release, distribution unlimited
Availability: Hard copy only.

This paper is part of the following report:

TITLE: KSCO 2002: Second International Conference on Knowledge
Systems for Coalition Operations

To order the complete compilation report, use: ADA402533

The component part is provided here to allow users access to individually authored sections
of proceedings, annals, symposia, etc. However, the component should be considered within
the context of the overall compilation report and not as a stand-alone technical report.

The following component part numbers comprise the compilation report:
ADP012330 thru ADP012354

Issues of Dynamic Coalition Formation Among Rational Agents

Matthias Klusch and Andreas Gerber

German Research Centre for Artificial Intelligence
Multiagent Systems Group
Stuhlsatzenhausweg 3, 66123 Saarbrücken, Germany
{klusch,agerber}@dfki.de

Abstract. We introduce the notion, issues, and challenges of dynamic coalition formation (DCF) among rational software agents in open, heterogeneous and world widely distributed environments such as the Internet and Web. Selected relevant approaches coping with only parts of the DCF problem domain in different disciplines such as decision theory, social reasoning, and machine learning are briefly discussed. Finally, we sketch one novel DCF scheme, and highlight some future research work towards a general framework of dynamic coalition formation.

1   Introduction

Self-interested, autonomous software agents on the Internet may negotiate rationally to gain and share benefits in stable (temporary) coalitions. This is to save costs by co-ordinating activities with other agents. For this purpose, each agent determines the utility of its actions and productions in a given environment by an individual utility function. The value of a coalition among agents is computed by a commonly known characteristic function which determines the guaranteed utility the coalition is able to obtain in any case. In a characteristic function game the agents may use imposed individual strategies to achieve a desired type of economically rational behaviour such as altruistic, bounded rational, or group rational. In any case, the distribution of the coalition's profit to its members is de-coupled from its obtainment but is supposed to ensure individual rational payoffs to provide a minimum of incentive to the agents to collaborate.

Rational agents should also be able to form beneficial coalitions in open, distributed and heterogeneous environments at any and in reasonable time. That includes scenarios in which dynamically occurring events may interfere with the running coalition processes such as continuous change of tasks to be accomplished, information and computing resources available to the agents, as well as temporary disconnection of coalition partners in the network, and changes in their reputation and trust.

Due to its nature dynamic coalition formation methods promise to be particularly well suited for applications of ubiquitous and mobile computing, including mobile commerce. M-commerce as it may be supported by personalised, rational information agents residing, for example, on WAP-enabled access devices such as pagers, organisers, (sub)notebooks, or UMTS cell phones, currently still remains to be an appealing vision for the common Internet user. However, the development and application of DCF methods enabling potential business partners to form temporary coalitions on demand, on the fly, at any time may inherently enable and even advance the development of effective mobile commerce and collaborative work. This includes, for example, the challenge of quickly forming time-constrained, profit-oriented customer coalitions for optimally negotiating, purchasing and sharing appropriately partitioned sets of items at multiple electronic market places world wide in reasonable time. First approaches into this direction include, for example, (Tsvetovat & Sycara, 2000; Lerman & Shehory, 2000; Preist, Byde & Bartolini, 2001; Yamamoto & Sycara, 2001; Shehory, 2001).

The remainder of this paper is structured as follows. Section 2 summarises some static approaches of forming stable coalitions among rational agents. Issues and problems of dynamic coalition environments are discussed in section 3 while selected relevant approaches to cope with parts of these problems are surveyed in section 4. We sketch a novel DCF scheme in section 5, and conclude the paper with a brief outlook on future work.

2   Static Formation of Stable Coalitions

According to (Conte and Sichman, 1995) models of coalition formation may be classified into two main approaches: utility-based and complementary-based models dividing the societies of actors into ones following either the principle of 'bellum omnium contra omnes' as it is largely favoured, for example, by game theory (Luce and Raiffa 1957, Axelrod 1984), or ones which rely on the collaborative use of complementary individual skills to enhance the power of each agent to accomplish its goals, respectively.

Up to now, most classic methods and protocols for a formation of stable coalitions among rational agents follow the utility-based approach. They rely on derived concepts from co-operative game theory, economics, and operations

research. Utilitarian coalition formation covers two main activities: (1) the generation of coalition structures, that is partitioning or covering the set of agents into coalitions, so as to maximise the monetary value depending on the benefit of accomplishing tasks regarding used resources and time spent; (2) the distribution of gained benefit among the participants of each of the coalitions. These activities may be interleaved and are not independent. A comprehensive discussion and classification of relevant work on coalition formation is given, for example, in (Kraus, 1997; Vauvert & El Fallah-Segrouhni, 2001).

2.1 Prerequisites

We briefly summarise the basic concepts and notions of co-operative game theory which are necessary to follow the discussion of coalition formation methods and the problems of dynamic coalition formation in subsequent sections. For a more comprehensive introduction to co-operative game theory we refer the reader to (Kahan & Rapoport, 1984; Osborne & Rubinstein, 1994; Holler & Illing, 2001).

2.1.1 Co-operative Games, Coalition Configurations

A **co-operative game** *(A, v)* is determined by a set *A* of agents wherein each subset of *A* is called a coalition, and a real-valued characteristic function *v: P(A)→R*, assigning each coalition its maximum gain, the expected total income of the coalition (the so-called coalition value). It is commonly assumed that (a) the value of any coalition *C* is in money, (b) the value *v(C)* does not depend on the actions of agents outside the coalition, (c) any coalition *C* forms by binding agreement on the distribution of its coalition value *v(C)* among its members, in particular no side-payments are allowed from *C* to any agents outside *C* within the game, and (d) the characteristic function *v* is known to all agents in *A*.

The solution of a co-operative game with side payments is a **coalition configuration** *(S,u)* which consists of
- a partition *S* of *A*, the so-called coalition structure, and
- an efficient payoff distribution $u : A \to \Re, (u(a_i))_{i \in \{1,...n\}} \in \Re^n, |A| = n$

The payoff distribution assigns each agent in *A* its utility out of the value of the coalition it is member of in a given coalition structure. It is commonly assumed that every coalition may form, including singletons or the grand coalition *A*. However, the number or size of coalitions to be formed using a coalition formation method is often restricted to ensure, for example, polynomial complexity of the formation process.

Individually rational distributions are assigning each agent at least the gain it may get without collaborating within any coalition, i.e., $\forall a \in A : u(a) \geq v(\{a\})$, it is assumed to hold for any coalition configuration. For group rational distributions it holds that $\forall C \subseteq A : \sum_{a \in C} u(a) \geq v(C)$,i.e., the group of all agents is assumed to maximise its joint payoff.

In coalition configurations with so-called Pareto-optimal payoff distributions no agent is better off in any other valid payoff distribution for the given game and coalition structure. A coalition configuration *(S,u)* is called stable if no agent has an incentive to leave its coalition in *S* due to its assigned payoff *u(a)*. Each notion of stability defines a particular solution space for co-operative games. Concepts of stability applied to coalition configurations are discussed in the context of coalition formation methods in the following section 2.2.

2.1.2 Coalition Algorithm, Coalition Formation Environment and Model

Rational agents which are involved in a co-operative game *(A,v)* are supposed to negotiate a stable payment configuration *(S,u)* as a solution of the game by the use of an appropriate **coalition algorithm** CA which should have the following desirable properties.
- *Local execution.* Each agent is able to execute the CA locally. Negotiation according to the CA is completely decentralised.
- *Anytime.* After any regular termination of an arbitrary co-operative game in the considered environment the CA outputs a stable configuration as a solution of that game.

A **coalition formation environment** *CE* for a given set of agents *A* is the set of assumptions and constraints which are valid for any kind of coalition forming activity between agents in *A* including propositions on
- The functionality of each of the agents in *A*, including, for example, the sets of tasks, actions, and utilities of its task-related productions,
- Valid methods for computing the values of coalitions, for example, by the sum of production utilities of all agents in a coalition,
- Valid methods for determining coalition configurations, including methods for searching coalition structures, negotiation and payoff distribution schemes.
- Commitments, obligations of and agreements between agents in *A* concerning the type of collaboration and interaction.

In a given coalition formation environment the agents particularly agree on (a) what kind of stable coalitions shall be negotiated (the considered notion of stability), and (b) what particular coalition algorithm CA shall be used for the negotiation. Please note that agents may, for example, use different utility functions to evaluate the utilities of task execution and corresponding productions.

A coalition environment is called *super-additive* or *sub-additive* depending on the type of all co-operative games it allows, and *general* if it allows for both, sub-additive and super-additive games. In non-super-additive environments at least one (all) pair(s) of potential coalitions is not better off by merging into one which could be caused by, for example, communication and co-ordination overhead costs, decrease of coalition value as a result of restricting utility constraints posed by agents joining a coalition, or anti-trust penalties for specific coalitions (Kraus & Shehory, 1999).

A **coalition formation model** CM = (CE, CA) is defined by both, the considered environment CE and given coalition algorithm CA for this environment. Interesting models are those where coalition formation is concerned with general and sub-additive environments. In environments where published interests and utilities used for negotiation to form coalitions cannot be verified, most current coalition algorithms allow for fraud by different types of lies. Arbitration schemes for competing agents with conflicting interests may help to circumvent such situations (Tesch and Fankhauser, 1999).

2.2   Selected Coalition Formation Methods

As mentioned above, current coalition formation methods aim at building stable coalitions. The meaning of stability of coalitions varies dependent from the considered application domain and discipline. Many if not most of the coalition formation algorithms today rely on chosen game-theoretic concepts for pay-off division within coalitions according to, for example, the Shapley-value, the Core, the Bargaining Set, or the Kernel (Kahan & Rapoport, 1984). We briefly discuss selected main approaches to (static) coalition formation based on co-operative game theory in subsequent sections[1].

2.2.1   Core-stable coalitions

One approach to form stable coalition configurations as proposed in (Sandholm, 1999) comprises the following two steps: Searching for a social welfare maximising coalition structure in a corresponding coalition structure graph for the given game (A, v), and then compute its payoff division according to the stability concept of the core (Wu, 1977). The core of a game with respect to a given coalition structure is the set of coalition configurations with not necessarily unique payoff distributions such that no subgroup of agents is motivated to depart from the given structure. Only coalition structures that maximise the social welfare, i.e., the sum of all coalition values of coalitions in the considered structure, are Core-stable. However, searching for an optimal coalition structure (given a set A of agents) among the exponential number of $|A|^{|A|/2}$ possible coalition structures is computationally hard since one has to try at least 2^{A-1} coalition structures (Sandholm et al., 1998). Another well-known problem with core-stable configurations is that the core may be empty for certain co-operative games, and is exponentially hard to compute. This hardly suits the needs of solution approaches for dynamic coalition formation.

2.2.2   Shapley-value stable coalitions

Any pay-off division scheme according to the so-called Shapley-value (Shapley, 1953) provides an agent the added value (marginal contribution) that it brings to the given coalition structure, averaged over all possible joining orders. Obviously, the Shapley-value is exponentially hard to compute. In contrast to the core the Shapley-value is proven to uniquely exist, to be Pareto-optimal, and individual and group rational for super-additive games.

Algorithms for forming stable coalitions which rely on the stability concept of the Shapley-value and a variation of it, the so-called bilateral Shapley-value (Ketchpel, 1994) applied to arbitrary n-agent co-operative games, are proposed in (Klusch, 1997; Klusch & Shehory, 1996b; Contreras et al., 1997). It is shown in (Klusch, 1997) that the computation of proposed payoff division according to the bilateral Shapley-value with equal or history-based recursive share among coalition members is of polynomial complexity, and is guaranteed to be efficient and individual rational for super-additive games. However, since it is also shown that the latter fact does not necessarily hold for sub-additive games, these algorithms are not suitable to dynamic environments in their current form. Ongoing research is performed to devise novel methods for adapting these algorithms to such environments.

2.2.3   Kernel-stable coalitions

The Kernel of a co-operative game *(A,v)* with respect to a given coalition structure is the set of so-called K-stable configurations *(S,u)* in which all coalitions in *S* are in equilibrium. Coalition *C* is in such an equilibrium if each pair of agents in *C* is in equilibrium, i.e., any pair of agents in *C* is balanced, that is, none of both agents can outweigh the other in (S,u) by having the option to get a better payoff in coalition(s) not in *S* excluding the opponent agent. In other words, agents argument each other like "Since I could obtain more without you in alternative coalitions than you without me, I deserve more, but without going to harm you." For this purpose each agent has to compare its surplus with those of other agents; the calculation of the surpluses bases on that of the excesses of all alternative coalitions. Obviously, the kernel of a game is exponentially hard to compute unless, for example, the size of the coalition is limited by a constant. The kernel appears to be attractive due to the following features: The kernel K is

---

[1] One publicly available simulation environment for coalition formation among rational information agents based on selected classic coalition theories is, for example, COALA (Klusch & Vielhak, 1997).

unique for any 3-agent game (A,v), assigns symmetric agents of some coalition in a given coalition structure for (A,v) equal payoff, and is locally Pareto-optimal in K.

Polynomial coalition algorithms for polynomial K-stable coalition configurations have been developed and applied to the domain of co-operative information agents in (Klusch & Shehory, 1996b; Shehory & Kraus, 1996b; Klusch, 1997).

2.2.4 Fuzzy coalitions

Negotiation during the coalition forming process may be connected with various forms of uncertainty. Such uncertainties could be induced by the possibility of dynamically occurring events which, for example, may hamper the negotiation process and produce vague or incomplete knowledge on expected profits or the share of the income of coalitions in which they intend to participate. This in turn implies so-called fuzzy co-operative games with vague profits and has been dealt with in numerous works, for example, in (Mares, 2001; Aubin, 1981). A fuzzy co-operative game with side payments is consisting of a set of agents and a fuzzy characteristic function v, and the membership function m of the fuzzy quantities $v(C)$ which may be interpreted as vague expectation of the common coalition profit that is to be distributed among its members. That is, the worth $v(C)$ of a coalition C is a fuzzy set of its (possible) real-valued coalitional profits. This set of fuzzy quantity $v(C)$ has at least one modal value, i.e., $m(v(C))=1$, determined by the membership function m. If for a given fuzzy co-operative game the coalition value $v(C)$ is equal to one modal value of C for all possible coalitions C, it is equivalent to a (deterministic) co-operative game. The vagueness of the distributed profit $v(C)$ means that particular payoff distributions can be realised with certain possibility only, which in turn is derived from the membership function m. Concepts of fuzzy super-additive co-operative games and "stable" fuzzy payoff distribution according to the fuzzy extension of the core and the Shapley-value are introduced and investigated in detail in (Mares, 2001). However, additional basic research on, for example, fuzzy sub-additive games and other concepts of "vague" stability remains to be performed, in particular appropriate coalition algorithms for fuzzy co-operative games have to be developed. This is topic of current research, for example, at DFKI.

2.2.5 Stochastic coalitions

Another class of co-operative games arises from co-operative decision making problems in stochastic environments. The notion of so-called stochastic co-operative games or co-operative games with stochastic payoffs, is introduced and investigated in (Suijs 1998; Suijs et al., 1999). A game with stochastic payoffs is defined by a set of agents, a set of possible actions coalitions may take, and a function assigning to each action of a coalition a real valued stochastic variable with finite expectation, representing the payoff to a coalition when this particular action is taken. Thus, in contrast to a deterministic co-operative game, the payoffs can be random variables, and the actions a coalition can choose from are explicitly modelled since the payoffs are not uniquely determined. It has been proven in (Suijs & Borm, 1999) that convex stochastic co-operative games are super-additive and have a non-empty core. Efficient coalition algorithms using these concepts are currently under development at DFKI.

However, all of the above mentioned as well as the vast majority of known other mechanisms for building utilitarian coalitions among agents remain static in the sense that they do not allow for any type of dynamic interference of running coalition formation processes. We will discuss types of dynamic events, corresponding problems and relevant approaches in the following sections.

3 Towards Dynamic Coalition Forming

The domain of *dynamic coalition formation* (DCF) among rational agents can be defined by the set of co-operation methods, schemes, and key enabling technologies to cope with the problem of dynamically building beneficial coalitions among agents in open, distributed, and heterogeneous environments such as the Internet.

3.1 The DCF Problem

The DCF problem rises in any collaboration environment and scenario in which at any time

(1) agents may enter or leave coalition formation processes,
(2) the set of tasks to be accomplished and the (computational) resources used, as well as
(3) the information, network, and user environment of each of the agents and the system as a whole may dynamically change.

Classical game-theoretic notions of coalition stability and respective negotiation algorithms are not applicable to such dynamic settings. Scenarios inducing uncertain, time-limited, context-based utilities and coalition values exacerbate the DCF problem. For example, an agent may determine the degree of membership to potential coalitions based on bargaining and the possible level of its commitment indicating the degree of collaboration that it desires.

3.2   Dynamic Coalition Formation Environments

As mentioned above, environments and settings in which rational agents have to be able to dynamically build coalitions can be characterised by the following classes of events and induced problems.

- *Tasks*: The set of tasks, goals and corresponding plans to accomplish may change for each individual agent at any time. Such changes concern, for example, the volume of tasks, utilities and costs of task execution as well as the frequency of such changes. This requires an agent to be able to perform, for example, fast dynamic re-planning of task execution to achieve its individual and/or common goals of the coalition. Re-planning concerns the granularity, re-usability and partiality or completeness of each of the considered plans. General task allocation problems are known as at least NP-hard problems. Real-time issues and requirements to perform planning under time-dependent uncertainty (Wellman, Ford & Larson, 1995) may even exacerbate these kinds of problems.
- *Agents*: Agents may leave or enter the agent society at any time, some agents may even temporarily hide their existence to parts of the society for different reasons.
-
- Optimisation:
- Negotiation:

We may distinguish between external and internal dynamic events. External events include, for example, a change in the specification of the problem to be solved by the agents, or any other change in the environment which are not caused by and cannot be influenced by the agents per se. Whereas internal events may be caused by the agents itself such as, for example, the entering or leaving of a coalition.

In dynamically changing environments rational agents may have to compute their individual utilities based on a pure sequence of local decisions. The problem of calculating an optimal complete mapping from states to actions (a so-called policy) in an accessible, stochastic environment with a known transition model is called a Markov decision problem. A transition model refers to a set of probabilities associated with the possible transition between states after any given action. Thus the agent is concerned with computing a sequence of values of stochastic variables X_t each of them is determined solely by the previous one. The resulting chain of probabilities $P(X_t|X_{t-1})$ yields a so-called Markov chain, a state evolution model. However, Markov chains and underlying decision support policies appear to be hardly feasible in open and dynamic environments for coalition formation. (Choi & Liu, 2001) propose one approach to mitigate the problem of prior knowledge on probabilities by using additional statistical information for the agents including the probability distributions of specific events to maximise their expected utilities without the need to of speculating others' actions. It remains to be investigated to what extent this approach can be generalised to coalition formation environments.

4   Selected Relevant Work

Relevant work on fuzzy coalition forming and co-operative games with stochastic payoffs (section 2.2), as well as rational revision of preferences, and other qualitative approaches to decision making based on partial, uncertain, and tentative information hold promise to be useful for coping with some of the issues of the DCF problem. We briefly discuss only some of the most relevant approaches and systems which are relevant for coping with parts of this problem. Other relevant work includes, for example, utility-based schemes for dynamically re-organising organisational structures (Barber & Martin, 2001), and exception tolerant reasoning and multi-criteria decision making under uncertainty (Benferhat et al., 2001; Dubois et al., 2000). These works may be properly extended for application to different dynamic coalition formation settings. The same hold with applying work on dynamic constraint satisfaction problems (Schiex & Verfaillie, 1993) since many of the above mentioned problems can be viewed naturally as CSPs (Eaton, Freuder & Wallace, 1998).

4.1   Game-Theory Based Approaches

4.1.1   Fuzzy and Stochastic Coalitions

Work on fuzzy and stochastic co-operative games as briefly described sections 2.2.4 and 2.2.5, respectively, is assumed to play an important role for the development of DCF schemes. Reasonable solutions for such types of games may lied to co-operation schemes which enable the agents to cope with issues of uncertainty, including, for example, vagueness of expected coalition values and corresponding payoffs. Such uncertainties may be induced by dynamic events such as network faults, changes of trust or reputation ratings of possible coalition partners, and receiving vague or even incomplete information and data during task execution or negotiation.

Both, the field of fuzzy and stochastic co-operative games still are in its very infancies and require further basic research efforts. This is even more valid for the application of principles and methods for such non-classical but still static coalition forming to dynamic settings. The development of algorithms for dynamic fuzzy or probabilistic

coalition forming appears to be most promising and challenging at the same time. We are currently working on the development of such DCF algorithms.

4.1.2 Overlapping Coalitions

A method for building overlapping coalitions for precedence-ordered task-execution has been proposed in (Shehory & Kraus, 1996c). The suggested any-time algorithm is of polynomial complexity and yields sub-optimal results. Goal satisfaction by agents is approached as a problem of assigning goals to coalitions of agents. Thus the distributed algorithm tries to compute appropriate partitions of the considered set of agents adopting solution methods (Chvatal, 1979) for the similar set covering problem which is known to be NP-complete (Cormen, Leierson & Rivest, 1990). The algorithm is relevant for dynamic environments, wherein the time period for negotiation and coalition formation may be changed during the process.

4.2 Social Reasoning

Social reasoning mechanisms are considered as essential building blocks suitable to situations where agents may dynamically enter or leave the society, without any global control. Such mechanisms are often based on the notion of social dependence (Castelfranchi et al., 1992), or aim at reputation and trust management.

4.2.1 Social Dependence Networks

In order to acquire and use dependence knowledge on the considered agent society each agent has to (a) explicitly represent some properties of the other agents, which may change dynamically, (b) exploit this representation thereby optimising its behaviour according to the evolution of the society, and (c) to monitor and revise its representation to avoid inconsistencies to an acceptable degree, without any pre-established global control.

For example, the multi-agent system DEPINT (Sichman, 1995) illustrates some essential aspects of an agent's social reasoning mechanism in particular concerning the (a) adaptation of an agent to changes in goals and plans, (b) formation of coalitions for plan achievement, and (c) revision of inconsistent belief. Each DEPINT agent dynamically builds and maintains its individual network of dependency relations with respect to the accomplishment of goals based on the skills of its own and that of other agents in the agent society[2]. It may adapt to changes in goals to pursue and corresponding feasibility of plans to perform by using this dependency knowledge to select at any moment the goals and plans which it actually is able to execute by itself and/or with the help of the society. The agent evaluates the susceptibility of other agents to adopt its goals which in turn enables it to dynamically form respective coalitions for accomplishing its tasks.

However, DEPINT agents are assumed (a) to show benevolent behaviour in the sense that they do not try to exploit each other, never offer erroneous information deliberately and always communicate information in which they believe; (b) posses complete and correct knowledge of their own goals, expertise, etc., and (c) to perform belief revision once inconsistent or contradictory belief about others is detected. These assumptions appear unrealistic in open, dynamic coalition environments as described above.

4.2.2 Reputation and Trust Management

Social mechanisms of reputation management aim at avoiding interaction with undesirable participants and may complement other security technologies for authentication and authorisation. Mechanisms for building, propagating, measuring and maintaining reputation and trust (Yu & Singh, 2000; Manchala, 2000) are useful to apply, for example, to settings for coalition formation among self-interested agents in e-commerce applications where trusted third parties are required but not available. Negotiation schemes for uncertain games with trusted third party are proposed, for example, in (Wu & Soo, 1999; Soo, 2000). The merging of several individual trust matrices which are commonly used as a means for assessing trust relationships is not necessarily transitive and certainly requires further research.

In general, mechanisms which allow agents to efficiently react on frequent changes of reputation ratings and assessment of trustworthiness of potential coalition partners with respect to, for example, the expected share of profits, reliability of membership, and benevolence are, to our knowledge, more than rare up to date. First approaches into this direction include, for example, fuzzy models of reputation in multi-agent systems (Rubiera, Lopez & Muro, 2001).

4.2.3 Time-Constrained Reasoning

Rational agents may face many potentially beneficial choices related to the timing of events which may occur during (a) the individual decision process, and/or (b) the negotiation process with other potential coalition partners.

---

[2] A DEPINT agent is said to be dependent on another if the latter may facilitate or prevent it from achieving one of its goals. Both agents are mutually or reciprocally dependent on each other with respect to the same or different goals, respectively

96

Regarding the use of social reasoning mechanisms in continuously changing environments temporal dependence networks and adequate temporal social reasoning mechanisms are proposed, for example, in (Allouche, Boissier & Sayettat, 2000). These mechanisms may be applied to DCF schemes which rely in part on social reasoning.

Relevant work on real-time issues in the context of agent-based online auctions (on a single auction server) suggesting a design for maximal asychrony and robustness to network delay includes, for example, (Wellman & Wurman, 2000). (Choi & Liu, 2001) propose a dynamic mechanism for simple but time-constrained trading. The preliminary results and experiences reported in these and other relevant work may be taken into account for a design of more complex dynamic customer coalition formation schemes.

5   One DCF Scheme: DCF-A

In this section we propose a DCF scheme, called DCF-A, to enable rational agents to react on events which occur dynamically during the coalition forming process. In this paper we do not focus on the details of the coalition forming according to some given coalition model but on the simulation of the

Due to the dynamic nature of the environment in which the agents are situated in their behaviour may change over time. We include appropriate learning components into the DCF scheme DCF-A to adapt the individual pay-off matrix of each agent to the current situation using reinforcement learning (Sutton & Barto, 1998), especially Q-learning (Mitchell, 1997). The main idea is to approximate the function assigning each state-action pair the highest possible pay-off. Regarding the adaptation of each agent's world model to frequent changes in the agent society we adopt the concept of levelled reasoning on the behaviour of other agents as it is described in (Weiss, 1999).

In the DCF-A scheme each coalition built is represented by one distinguished agent acting as the so-called coalition leader. The coalition leader continuously attempts to improve the value of its coalition. In order to prevent the implied communication overhead between the leader and other members of the coalition, the leader simulates possible adjustments of the actual coalition configuration by building hypothetical re-configurations and rating them based on the members' capabilities, resources, desirability, communication stability, task description, and suggestibility from the current environment. As soon as the coalition leader achieves a significant improvement of the coalition value by simulation, it informs all its coalition members about proper alternatives. In turn, the agents have to send their estimation about the quality of relevant services and agents in regular time periods to the coalition leader or some so-called world utility agents. This is quite similar to the co-ordination and collaboration within so-called holonic multi-agent systems (Gerber, Siekmann & Vierke, 1999).

The coalition leader is assumed to be able to obtain up to date information about the agent society, for example, by request from some distinguished so-called 'world utility agent'. Such world utility information include public rankings about the quality of services offered by individual agents. Each agent may get a vague idea of the utilities and estimated payoffs of other agents, services, etc. When a new agent initialises itself and has no or less information on the world's entities, a global world utility function can give him a first hint while deciding what is a good choice to do next. The world utility on the one hand (in a benevolent agent society) can be used to give a global guideline for later evolution of the society. On the other hand (in a non-benevolent society) a group of agents may try to manipulation the world utility of some items for their own interests. But as more agents report their own estimation about entities listed at the world utility agent, the harder it will be to manipulate these utilities. Therefore we extend the world utility function by collecting the number of remarks from different agents for one ranked entity. Only the newest remark from an agent about an entity is stored. In addition, to avoid the world utility value from jumping from low to high, we extend the world utility function with proper learning mechanism. The world utility function provides a median of the incoming remarks and may provide common utility estimations of relevant items, entities and relationships of the society.

The DCF-A Scheme (Dynamic Coalition Formation Based on Simulation)

Variables and functions used by the DCF-A:
C configuration of a coalition (members, payoffs)
CPL list containing the changes (new partners) in of the coalition structure in relation to the current structure
AAL list containing the agents' abilities (capabilities, capacity, desirability, communication stability, stability of task description, suggestibility from the environment)
tp trust penalty for removing an agent from coalition C
cv current value of coalition C based on the Shapley-value
$rvf()$ function to determine the risk value when adding an agent a_i to coalition C (Linsmeier & Pearson, 1996; Alexander, 1998)
Individual agent's preferences characterising its behaviour:
wr worst acceptable risk to remove a single agent a_i from C and getting punished from the agent society by loosing reputation

wtp worst acceptable trust penalty for which the coalition head is willing to change the coalition structure with regards to all agents of the current *CPL*

k number of simulation cycles as an upper bound for the number of agents that have to be requested during the negotiation phase ($|C| \leq k$). A higher *k*-value denotes a higher risk in not getting all the changes of the coalition structure realised, but the chance to obtain a higher performance of the coalition is also higher.

Coalition formation and adjusting protocol used by each of the coalition leaders:

1. Initialisation Phase
   CPL = null
   halt = false

2. Simulation Phase

To prevent to get stuck in a local maximum and to avoid cyclic changes of the coalition structure, we use a randomised version of the algorithm for the simulation phase. The algorithm for the simulation phase is intended to run as long as it is not necessary to make changes of the coalition structure. In case of the occurrence of dynamic events it stops and presents a valid configuration which does not decrease the coalition value compared to that in the previous configuration. Therefore the agent does not change the current configuration, instead it builds hypothetical coalition structures and configurations, and simulates possible changes of them. During these iterations the actually best solution is stored in *BestCPL* such that the algorithm can be halted at any time and outputs a valid solution. The solution is not a degeneration of a previous solution since the simulation phase is stopped if and only if the value of the hypothetical configuration appears to be much better then that of the current configuration. The argument 'much better' is necessary to prevent too many changes in the coalition structure. The simulation phase is an any time algorithm.

   while not (*halt*) **do**
      requesting *newAAL* from distinguished world utility agent
      merging *newAAL* with local *AAL*: For this purpose we adopt learning mechanisms (Watkins, 1989; Sutton & Barto, 1998) and stochastic methods for agent ratings;
      CPL := null
      for (*c*=1 to k)
         choose randomly one operation for cycle *c* (noop, add_member, remove_member)
         if add_member **then**

            choose agent a_i from *AAL* with [$min_{1 \leq i \leq |AAL|} rvf(a_i)$ and $max_{1 \leq i \leq |AAL|} value\{C+a_i\}$]
            insert tupel [a_i , add] to *CPL*
         if remove_member **then**

            choose agent a_i from *AAL* with [$max_{1 \leq i \leq |C|} rvf(a_i)$ and $max_{1 \leq i \leq |AAL|} value\{C-a_i\}$]
            if $rvf(a_i) > wr$ **then**
               insert tupel [a_i , remove] to *CPL*
               $tp := tp + 1/rvf(a_i)$
      next
      if value(*CPL*) > value (*LastCPL*) **then**
         // following types of dynamic events are considered: changes of the current coalition configuration, or changes in the environment or task requirements.
         BestCPL=CPL
      If value(*BestCPL*) >> *cv* and *tp<wtp* **then**
         // if a new coalition structure is found that is much better then the old one, then the simulation is stopped and the negotiation phase for realising the hypothetical coalition re-configuration begins
         halt = true
   while end

3. Negotiation Phase

Concerning the fact of a dynamic environment the term of stability of a coalition has to be properly modified. In our case of a dynamic scenario it is not possible to build stable coalitions in the classical game-theoretic sense. This is because at any time dynamic events may happen and the coalition configuration has to be adjusted in real-time. However, in situation where no dynamic events occur, the rankings of the agents are stable, the simulated coalition protocol finds the approximately best configuration (if it exists) and hold it until a change in the environment happens. After the simulation phase has stopped the *BestCPL* is used in the following negotiation phase, where the coalition leader tries to realise the corresponding hypothetically "best" configuration. It sequentially gets into a

negotiation process with each agent of the *BestCPL* list based on a mechanism for 'multi-attribute negotiations' (Jonker & Treur, 2001). The agents have to negotiate about multiple attribute values, for example, the remaining time to fulfil a particular service, the costs of the service, etc. It is not guaranteed that all negotiations will end successfully. Thus, we adopt a 'levelled commitment protocol' (Andersson & Sandholm, 2001).

> *halt := false*
> **for** (*i=1 to |BestCPL|*)
>   [a_i, operation$_i$]:= *i*-th tupel of *BestCPL*
>   **try**
>   **if** operation$_i$ = add_member **then**
>   bilateral negotiation with agent a_i based on protocols for multi-attribute negotiation and 'levelled commitment contracts' [1] (if not all agents of the BestCPL can be added to this coalition).
>   **if** negotiation was successful **then**
>   **add** a_i to C
>   **else**
>   **remove** a_i from C
>   **catch** (if any dynamic event occurs during the execution of the negotiation phase)
>   **stop** Negotiation Phase
> **next**

4. Evaluation Phase

Send *AAL* to the known world utility agent, which merges this list with its local *AAL* (using learning mechanisms and stochastic methods for the agent rankings). Restart the simulation phase (Go to 2.)

6   Conclusions

We introduced the notion, selected issues, and challenges of dynamic coalition formation (DCF) among rational software agents. In addition, we briefly discussed selected relevant work in different disciplines and proposed a novel DCF scheme. It has to be emphasised that one of the main challenges of the domain of dynamic coalition formation is the development of efficient DCF algorithms which enable rational agents to efficiently cope with different hard issues and problems they are facing in continuously changing, open, distributed and heterogeneous environments such as the Internet and Web. This is one focus of ongoing and future research, for example, at DFKI. For this purpose, many relevant approaches and theoretical work stemming from different disciplines are available to date including work on temporal social reasoning, and fuzzy and stochastic co-operative games.

References

Alexander, C. (1998) "Risk Management and Analysis, Risk Measurement and Management" ISBN: 0-471-97957-0, 1998.

Andersson, M. and Sandholm, T. (2001) "Leveled Commitment Contracts with Myopic and Strategic Agents" Journal of Economic Dynamics and Control, Vol. 25, pp. 615-640, Special issue on Agent-Based Computational Economics, 2001.

Aubin, J.P. (1981) "Cooperative fuzzy games" Journal Mathematics of Operations Research, Vol. 6(1), February 1981

Axelrod, R. (1984) "The Evolution of Cooperation" Basic Books, New York, 1984.

Barber, K.S. and Martin, C.E. (2001) "Dynamic re-organisation of decision-making groups" Proceedings 5[th] International Conference on Autonomous Agents, Montreal, Canada, ACM Press, May/June, 2001.

Benferhat,S., Dubois, D., Prade,H., and Sabbadin,R (2001) "Exception –tolerant reasoning and decision uner uncertainty using qualitative possibility theory" Journal Künstliche Intelligenz, Vol. 3, arendtap Verlag, Beremen, 2001.

Castelfranchi, C., and Micelli, M., and Cesta, A. (1992) "Dependence Relations Among Autonomous Agents" W Eric and Y Demazeau (eds.), Decentralized A. I. 3, Elsevier Science Publishers, pp 215-227, 1992.

Choi, S.P.M. and Liu, J. (2001) "A dynamic mechanism for time-constrained trading" Proceedings 5[th] International Conference on Autonomous Agents, Montreal, Canada, ACM Press, May/June, 2001.

Chvatal, V. (1979) "A greedy heuristic for the set-covering problem" Journal of Mathematics of Operations Research, Vol. 4(3), pp. 233-235, 1979.

Conte, R., and Sichman, J.S. (1995) "DEPNET: How to benefit from social dependence" Journal of Mathematical Sociology, 20(2-3), 161-177, 1995.

Contreras, J., Wu, F., Klusch, M. and Shehory, O. (1997) "Coalition formation in a power transmission planning environment", Proceedings of International Conference on Practical Applications of autonomous Agents and Multi-Agent Systems (PAAM-07), London, UK, April 1997.

Cormen, T.H., Leierson, C.E., and Rivest, R.L. (1990) "Introduction to algorithms" MIT Press, CA, USA, 1990.

Dubois, D., Grabisch, Modave, F., and Prade, H. (2000) "Relating decision under uncertainty and multicriteria decision making models" Journal of Intelligent Systems, Vol. 15, pp. 967-979, 2000.

Eaton, P.S., Freuder, E.C., and Wallace, R.J. (1998) "Constraints and agents – Confronting ignorance" AI Magazine, Vol. 19(2), 1998.

Gerber, C., Siekmann, J., and Vierke, G. (1999) „Holonic multi-agent systems". Technical Report Research Report RR-99-03, DFKI, 1999.

Holler, M.J. and Illing, G. (2000) "Einführung in die Spieltheorie" Springer Heidelberg, 4[th] edition, 2000

Jonker, C.M. and Treur, J. (2001) "An Agent Architecture for Multi-Attribute Negotiation" Proceeding of 17[th] International Joint Conference on Artificial Intelligence (IJCAI'01), Vol.2, pp. 1195-1201, 2001.

Kahan, J.P., and Rapoport, A. (1984) "Theories of coalition formation" Lawrence Erlbaum Associates, New Jersey.

Ketchpel, S. (1993) "Coalition formation among autonomous agents" Proceedings of European Workshop on Modelling Autonomous Agents in a Multi-Agent World (MAAMAW-93), 1993.

Ketchpel, S. (1994) "Forming coalitions in the face of uncertain rewards" Proceedings of AAAI Conference, AAAI-94, pp. 414-419

Klusch, M. (1997) "Kooperative Informationsagenten im Internet", Dissertation, Schriftenreihe Forschungsergebnisse zur Informatik, Dr. Kovac Verlag, Hamburg, ISBN 3-86064-746-6, 1997.

Klusch, M. and Shehory, O. (1996) "Coalition formation among rational information agents" Proceedings of European Workshop on Modelling Autonomous Agents in a Multi-Agent World (MAAMAW-96), W Van de Velde and JW Perram (eds.), Lecture Notes in Artificial Intelligence, Vol. 1038, pp. 204-217, 1996.

Klusch, M. and Shehory, O. (1996b) "A polynomial kernel-oriented coalition formation algorithm for rational information agents" Proceedings of International Conference on Multi-Agent Systems (ICMAS-96), Kyoto, Japan, December 1996, pp. 157-164.

Klusch, M. and Vielhak, T. (1997) COALA - Simulation Environment for Coalition Formation Among Autonomous Information Agents. http://www.dfki.de/~klusch/COALA/

Kraus, S. (1997) "Negotiation and cooperation in multi-agent environments", Journal of Artificial Intelligence, Vol. 94(1-2), pp. 79-98.

Larson, K. and Sandholm, T. (2000) "Anytime Coalition Structure Generation: An Average Case Study" Journal of Experimental and Theoretical AI, Vol. 12, pp. 23-42.

Linsmeier,T.J. and Pearson, N.D. (1996) "Risk measurement: an introduction to value at risk" ACE OFOR from University of Illinois at Urbana-Champaign. http://agecon.lib.umn.edu/uiuc/aceo9604.pdf

Luce, R. D. and Raiffa, H. (1957) "Games and Decisions: Introduction and Critical Survey", John Wiley & Sons Ltd., 1957.

Manchala, D.W. (2000) "E-commerce trust metrics and models" IEEE Internet Computing, Vol.4(2), March/April, 2000.

Mares, M. (2001) "Fuzzy cooperative games – Cooperation with vague expectations" Physica Verlag, 2001.

Rubiera, J.C., Lopez, J.M.M., and Muro, J.D. (2001) "A fuzzy model of reputation in multi-agent systems" Proceedings 5[th] International Conference on Autonomous Agents, Montreal, Canada, ACM Press, May/June, 2001.

Sandholm, T. (1999) "Distributed Rational Decision Making" G. Weiss (ed.), "Multiagent Systems: A Modern Introduction to Distributed Artificial Intelligence", MIT Press, pp. 201-258, 1999.

Sandholm, T. and Lesser, V. (1997) "Coalitions among Computationally Bounded Agents" Journal of Artificial Intelligence, 94(1), Special issue on Economic Principles of Multiagent Systems, pp. 99-137.

Sandholm, T., Larson, K., Andersson, M., Shehory, O., and Tohmé, F. (1999) "Coalition Structure Generation with Worst Case Guarantees" Artificial Intelligence, 111(1-2), pp. 209-238.

Schiex,T. and Verfaillie,G. (1993) "No-good recording for static and dynamic constraint-satisfaction problems" Proceedings of International Conference on Tools with AI, pp. 48-55, Washington DC, IEEE Computer Society Press, 1993.

Shehory, O. and Kraus, S. (1996a) "Cooperative goal-satisfaction without communication in large-scale agent-systems" Proceedings of European Conference on AI (ECAI-96), pp. 544-548, Budapest, Hungary, August 1996.

Shehory, O. and Kraus, S. (1996b) "A Kernel-oriented model for autonomous-agent coalition-formation in general environments: implementation and results" Proceedings of AAAI Conference, pp. 134-140, Portland, Oregon, August 1996.

Shehory, O. and Kraus, S. (1996c) "Formation of overlapping coalitions for precedence-ordered task-execution among autonomous agents" Proceedings of International Conference on Multiagent Systems (ICMAS-96), pp 330-337, Kyoto, Japan, December 1996.

Shehory, O. and Kraus, S. (1998) "Methods for task allocation via agent coalition formation", Journal of Artificial Intelligence, Vol. 101 (1-2), pp. 165-200, May 1998.

Shehory, O. (2001) "Optimality and Risk in Purchase at Multiple Auctions" Proceedings of 5[th] International Workshop on Cooperative Information Agents (CIA 2001), M Klusch, F Zambonelli (eds.), Lecture Notes in Artificial Intelligence, Vol. 2182, Springer, 2001.

Sichman, J.S. (1998) "DEPINT: Dependence-Based Coalition Formation in an Open Multi-Agent Scenario" Journal of Artificial Societies and Social Simulation vol. 1, no. 2, March 1998.

Soo, V.W. (2000) "Agent negotiation in trusted third party mediated uncertain games" Proceedings of 4[th] International Conference on Autonomous Agents, Barcelona, June, ACM Press, 2000.

Suijs, J. (1998) "Cooperative Decision Making in a Stochastic Environment" CentER Dissertation Series #36, CentER. Tilburg, 1998.

Suijs, J. and Borm, P. (1999) "Stochastic cooperative games: superadditivity, convexity and certainty equivalents" Journal of Games and Economic Behavior, Vol. 27, pp. 331-345, 1999.

Suijs, J., Borm, P., and De Waegenaere, A., and Tijs, S. (1999) "Cooperative games with stochastic payoffs" European Journal of Operational Research, Vol. 113, pp. 193-205, 1999.

Sutton, R.S., and Barto, A.G. (1998) "Reinforcement Learning: An Introduction", MIT Press, Cambridge, MA, Bradford Book, 1998.

Tohmé, F. and Sandholm, T. (1999) "Coalition Formation Processes with Belief Revision among Bounded Rational Self-Interested Agents" Journal of Logic and Computation, Vol. 9(6), pp. 793-815.

Tsvetovat, M. and Sycara, K. (2000) "Customer coalitions in the electronic marketplace" Proceedings of 4[th] International Conference on Autonomous Agents, Barcelona, June, ACM Press, 2000.

Vauvert, G., and El Fallah-Segrouhni, A. (2001) "Coalition formation among strong autonomous and weak rational agents", Proceedings of the 10[th] European Workshop on Modelling Autonomous Agents in a Multi-Agent World (MAAMAW), Annecy, France, May 2001.

Vauvert, G., and El Fallah-Seghrouchni , A. (2000) "Coalition formation among egoistic agents". Proceedings of International ICSC Symposium on Multi-Agents and Mobile Agents in Virtual Organizations and E-Commerce (MAMA'2000.), Wollongong, Australie, December 2000.

Watkins, C.J.C.H. (1989) "Learning from Delayed Rewards". PhD thesis, King's Collage. Cambridge University, 1989.

Wellman, M.P. and Wurman, P.R. (1998) "Real time issues for Internet auctions" Proceedings 1[st] Workshop on Dependable and Real-Time E-Commerce Systems (DARE-98), Denver, CO, USA, June 1998.

Wellman,M.P., Ford, M., and Larson, K. (1995) "Path planning under time-dependent uncertainty" Proceedings of 11[th] Conference on Uncertainty in AI (UAI-95), Montreal, Canada, August, 1995

Wu, L.S. (1977) "A dynamic theory for the classes of games with non-empty cores" SIAM Journal of Applied Mathematics, Vol. 32, pp. 328-338, 1977.

Wu, S.H. and Soo,V.W. (1999) "Risk control in multi-agent coordination by negotiation with a trusted third-party" Proceedings of International Joint Conference on AI (IJCAI-99), Stockholm, Sweden, 1999.

Yamamoto, J. and Sycara, K. (2001) "A stable and efficient buyer coalition formation scheme for e-marketplaces" Proceedings 5[th] International Conference on Autonomous Agents, Montreal, Canada, ACM Press, May/June, 2001.

Yu, B., Singh, M.P. (2000) "A social mechanism of reputation management in electronic communities" Proceedings of International Workshop on Cooperative Information Agents IV, M Klusch & L Kerschberg (eds.), Lecture Notes in Artificial Intelligence, Vol. 1860, Springer, July 2000